

EV316936945

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Communication Protocol for Synchronizing
Animation Systems**

Inventors:

Leonardo Blanco

Andrei Baioura

Matt Calkins

Paul David

ATTORNEY'S DOCKET NO. MS1-1756US

TECHNICAL FIELD

The systems and methods described herein relate generally to animation systems and, more specifically, to a communication protocol for maintaining data and clock synchronization between multiple animation systems.

BACKGROUND

Whether related to video games, network web sites, multimedia presentations or the like, user expectations with regard to graphics quality has significantly increased in recent years. Consequently, this high level of expectations poses ever-increasing challenges to video/animation system and application developers.

Interactive animation applications present a special consideration because two animation systems (e.g., an application stored in local or remote memory and a display device system) must communication with each other to provide a quality user experience. A problem exists here because one animation system (i.e. the application stored in memory) is mostly concerned with user interactivity while the other animation system (i.e. the display device system) is primarily concerned with rendering aspects such as refresh rate. These competing interests must be reconciled to provide the desired user experience.

SUMMARY

Systems and methods are described that include two distinct animation systems. A high level animation system (e.g., an application) is primarily concerned with interactivity and complex timing structure and, therefore, is optimized for high interactivity. A low level animation system (e.g., a display device) is primarily concerned with a rendering refresh rate and is thus optimized for a high refresh frame rate. The provision of multiple distinct animation systems allows animation to run at a predefined display refresh rate without penalizing interactivity, or vice-versa.

The two animation systems run asynchronously so that each system can better focus on its primary function in its own timing structure. To optimize the user experience, the two animation systems must be synchronized. In the synchronization process described herein, the high-level animation system and the low-level animation system exchange data (e.g., clock data, animation function data, etc.) with the use of a communication protocol designed specifically for this purpose.

The communication protocol introduced herein provides an efficient way to exchange the required data. Instead of sending a large amount of data for each frame, the communication protocol provides for sending just a small amount of data about how the animation is supposed to change over a specified period of time. In addition to saving system overhead, the protocol ensures that the low-level animation system has information to process several frames of an animation, which results in no rendering frames being dropped due to a lack of refresh data.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of exemplary methods and arrangements of the present invention may be had by reference to the following detailed description when taken in conjunction with the accompanying drawings wherein:

Fig. 1 is a block diagram of a server-client system including animation subsystems.

Fig. 2 is a block diagram of an exemplary animation rendering system architecture.

Fig. 3 is a flow diagram depicting a methodological implementation of a basic operation of a factorized animation/rendering system.

Fig. 4a is a depiction of communication protocol messages and parameters transmitted from a high level timing engine to a low level timing engine to synchronize the high level animation system with the low level animation system.

Fig. 4b is a depiction of communication protocol messages and parameters transmitted from a low level timing engine to a high level timing engine to synchronize a high level animation system with a low level animation system.

Fig. 4c is a depiction of communication protocol messages and parameters transmitted from high level animation objects to low level animation objects to synchronize a high level animation system with a low level animation system.

Fig. 5 is a flow diagram depicting an exemplary methodological implementation of an animation process in a rendering system having a fully connected high level animation system and low level animation system.

Fig. 6 is a flow diagram depicting an exemplary methodological implementation of an animation process in a rendering system having a partially connected high level animation system and low level animation system.

1 Fig. 7 is an exemplary computing environment in accordance with the
2 systems and methods described herein.

3 4 **DETAILED DESCRIPTION**

5 The following discussion deals generally with a computer system
6 generating and rendering an animation (i.e. a function that modifies graphical
7 objects over time) using a communication protocol that governs transmissions
8 utilized in the generating and rendering. The graphical objects may comprise a
9 window, a dialog box or the like. As used herein, the term “computer system” can
10 mean either a single computer or multiple computers working together. The
11 computer system includes multiple distinct animation/rendering systems - or
12 subsystems - that each focuses on a particular animation or rendering function.

13 In the discussion below, a communication protocol is described in the
14 context of two animation systems or subsystems. It is noted, however, that these
15 are only examples and that the communication protocol may be used with more
16 than two distinct animation systems or subsystems.

17 The term “rendering” is used in the discussion of the communication
18 protocol introduced herein. “Rendering” may be viewed as a step in an animation
19 process that defines the animation for display and/or displays the animation. A
20 “rendering system” can be a system that includes a rendering function. This may
21 be narrowly defined to describe a system that only performs a rendering function,
22 or it may be broadly defined to describe a system that performs other functions in
23 addition to one or more rendering functions, such as an animation system.

24 Similarly, the term “animation system” as used herein refers to a system
25 that includes an animation function. An “animation system” may be a system that

1 only performs animation functions, or it may be a system that performs one or
2 more animation functions in addition to other functions, such as a rendering
3 function.

4 The use of either particular term - "rendering system" or "animation
5 system" - is not meant to limit the scope of the architecture and/or function
6 associated therewith. The proper scope of the term used may be derived from the
7 context in which the term is used.

8 A full-featured animation system includes complex timing structures,
9 wherein each animation can be driven by a separate clock, and the clocks are
10 related by an arbitrary set of rules specified by an application developer. For
11 example, one animation (A) can be specified to begin exactly at the same time that
12 a different animation (B) ends, or a group of animations can be specified to run at
13 half their normal speed. In a more general sense, such a system can also handle
14 synchronization of animations with other types of linear media, such as audio and
15 video, by associating a clock with each piece of linear media.

16 In addition, the clocks are interactive, in that they can be started, paused or
17 stopped at any time, and those interactions may cause changes that propagate
18 through the entire timing hierarchy as dictated by the synchronization rules.
19 Following the previous example, if animation (B) is paused and resumed five
20 seconds later, then animation (A) will begin five seconds later than originally
21 planned. This, in turn, may cause timing changes in other clocks, potentially
22 propagating through the entire timing structure.

23 Calculating the changes to the timing structure is an expensive unbounded
24 operation, so this operation is not suitable for a real-time rendering system that has
25 a goal of meeting a predictable frame rate. These two goals (providing a full-

1 featured animation system and maintaining a predictable frame rate) are
2 contradictory. One solution is to divide the system into two parts: one which
3 knows about synchronization rules between clocks and another that considers each
4 clock as a fully independent entity.

5 The communication protocol claimed herein maintains the two distinct
6 animation systems in synchronization so that animations can run at a high display
7 refresh rate without penalizing interactivity and vice-versa.

8 An animation system described herein includes at least two components,
9 namely, a high-level animation subsystem that is primarily concerned with
10 interactivity and complex timing structures, and a low-level animation subsystem
11 that is primarily concerned with rendering one or more animations on a display at
12 a high, constant display refresh rate. The term "high-level" refers to the fact that
13 this animation system is closer to the user, while the "low-level" animation system
14 is closer to the metal (or machine).

15 One example of a high-level animation subsystem is an application that
16 controls animations according to data received from the application itself or from a
17 user. An example of a low-level animation subsystem is a display device system
18 that controls converting animation data to physically displayed animations. The
19 display device system includes hardware components as well as software
20 components that may be included on video card hardware (not explicitly shown)
21 or in a component that does not physically reside on the video card.

22 In the present discussion, reference may be made to a display device as
23 being a display system or an animation system or subsystem. However, it is noted
24 that all the components and processing of a display system are not necessarily
25 physically resident on the display device. The term "display device" as used

1 herein also encompasses processor operations that control a display and any kind
2 of display memory.

3 Normally, an animation process is thought of as being run in process on a
4 single machine that includes both the high-level animation subsystem and the low-
5 level animation subsystem (using a single animation system). However, the
6 communication protocol described herein for two distinct animation systems
7 anticipates the high-level and low-level animation subsystems running: (1) as two
8 threads in a single process, (2) as two distinct processes on a single machine; or
9 (3) on two distinct machines separated by a network, such as in a client-server
10 architecture.

11 **Exemplary Server-Client System**

12 **Fig. 1** is block diagram of an exemplary server-client system 100 in
13 accordance with the systems and methods described herein. The server-client
14 system 100 includes a server 102 and a client 104 that communicate over a
15 network 106, such as the Internet.

16 The server 102 includes memory 108, a processor 110, a network interface
17 card 112 for controlling communications over the network 106, and other
18 miscellaneous hardware 114 typically required for a server to perform standard
19 server functions. The memory 110 includes an operating system 116 and
20 miscellaneous software components 118 that provide server functionality.

21 An application 120 is also stored in the memory 108 and includes a high-
22 level animation subsystem 122 that incorporates computer code that controls one
23 or more animation functions. The animation function could be a window, a dialog
24 box, or the like. The high-level animation subsystem 122 is primarily focused on
25 interactivity and complex timing structures. As will be seen in further discussion,

1 the application 120 may include user-written computer code and system computer
2 code that is isolated from user code that the user code accesses via one or more
3 application programming interfaces (API).

4 The client 104 includes memory 130, a processor 132, a network interface
5 card 134 that controls communications to and from the network 106, and a display
6 136. The client 104 also includes a display device 138 as well as miscellaneous
7 hardware 140 that may be required for the client 104 to function appropriately.

8 The memory 130 stores a browser 142 configured to provide access to and
9 navigate the network 106 and an operating system 144 that controls the basic
10 functionality of the client 104. Miscellaneous software components 146 are also
11 stored in the memory 130 and include software that provides necessary client 104
12 functionality and client services.

13 The display device 138 includes a low-level animation subsystem 148 that
14 may be stored in display device memory (not shown), in a display device hardware
15 component (not shown) or in the memory 130 included in the client 104. Fig. 1
16 shows that at least a portion of the display device 138 components may reside in
17 the main memory 130 of the client 104. The low-level animation subsystem 148 is
18 primarily focused on meeting rendering targets for a high refresh frame rate (a
19 minimum refresh rate of thirty (30) frames per second). Sustaining the high
20 refresh frame rate provides an excellent result and is pleasing to a viewer, which is
21 the ultimate goal of an animation system.

22 Two obstacles that stand in the way of sustaining a high refresh frame rate
23 for an animation are (1) latency, and (2) bandwidth. When sending data over a
24 network, the data can be destined for anywhere in the world, which can create
25 significant latencies. A latency problem prevents appropriate messages and

1 feedback from arriving on time and this can cause the refresh rate to suffer and
2 degrade the user experience. This is a greater problem with animation systems
3 configured as single entity, since a great amount of data has to be sent regularly to
4 control the animation.

5 The bandwidth issue also presents a greater problem for single entity
6 systems. With a complex animation, huge amounts of data must be sent across the
7 particular boundary (thread, process, network) for every frame. Even with the
8 system described herein that includes two distinct animation subsystems, using
9 bandwidth can be a problem with the use of a modem that inherently limits
10 bandwidth or with a server that is serving an enormous number of clients. Even if
11 a server must only transmit a relatively small amount of data for a simple
12 animation, if that data must be served to, say, four hundred clients, it can become a
13 bandwidth issue.

14 The systems and methods described herein provide an efficient way to
15 minimize bandwidth utilization while maintaining a high refresh rate.

16 **Exemplary Animation Rendering System Architecture**

17 **Fig. 2** is a block diagram of an exemplary animation rendering system
18 architecture 200 in accordance with the systems and methods described herein.
19 Some of the elements shown in Fig. 2 are elements that are also included in Fig. 1.
20 For reference purposes, an element included in both figures is shown in Fig. 2 with
21 the same reference numeral used for the element in Fig. 1.

22 The exemplary architecture 200 includes the application 120 and the
23 display device 138. The application 120 communicates with the display device
24 138 via communication channel 204 according to a communications protocol 216
25 that will be discussed in greater detail below. As previously discussed, the

1 communication channel 204 traverses a boundary that may be a boundary between
2 threads, processes or machines. For a boundary between machines, for example,
3 the communication channel 204 may be the network 106 previously shown.

4 User code 202 functions together with the application 120 via an API
5 boundary 206 that is a set of system APIs that provide access to the application
6 120 by a developer that wants to utilize application 128 features with the user code
7 202.

8 As previously discussed, the application 120 encompasses the high-level
9 animation subsystem 122. The high-level animation subsystem 122 includes a
10 high-level graphics component 208, a high-level animation objects database 210
11 and a high-level timing engine 212. The timing engine 212 creates and controls
12 one or more high-level clocks 214 stored in the high-level animation object
13 database 210.

14 The high-level timing engine 212 is responsible for setting up and
15 synchronizing the high-level clocks 214. So, for example, if there are ten (10)
16 animations that are supposed to run together, the high-level timing engine 212
17 synchronizes ten (10) high-level clocks 214 to display the animations in
18 synchronicity.

19 The high-level graphics component 208 stores the type of graphic used in
20 an animation. The high-level animation objects database 210 stores animation
21 values associated with one or more animations (and/or media objects associated
22 with one or more media). The animation values include at least a timeline (i.e., a
23 clock) associated with the animation. (It is noted that the terms “clock” and
24 “timeline” as used herein are interchangeable). The animation values affect the
25

1 output of animation rendering based on values of the high-level clocks 214. Each
2 animation modifies one or more aspects of the high-level graphics component 208.

3 For example, for a simple animation of creating a line from point P1 to
4 point P2, the animation objects database 210 would store a value for P1 at time=0
5 and a value for P2 at time=1. (All points between P1 and P2 are interpolated).

6 The display device 138 of the exemplary animation rendering system
7 architecture 200 includes the low-level animation subsystem 148 shown in Fig. 1.
8 The low-level animation subsystem 148 is an equivalent structure to the high-level
9 animation subsystem 122 and includes a low-level graphics component 218, a
10 low-level animation objects database 220 and a low-level timing engine 222. The
11 timing engine 222 creates one or more low-level clocks 224 that are stored in the
12 low-level animation objects database.

13 The components (218 - 224) of the low-level animation subsystem 148 are
14 similar to the components (208 - 214) of the high-level animation subsystem 122
15 and their functions and inter-operability are similar as previously described.

16 There are advantages of an architecture like the exemplary animation
17 rendering system architecture 200, i.e. an architecture that has an animation
18 system factorized into two subsystems. If an application is expensive and spends a
19 lot of overhead with interactivity, the animation can continue; it doesn't have to
20 wait for the interactive portion to complete before the animation can continue. In
21 a non-factorized system, if the application doesn't provide the display with regular
22 data updates, then the display may draw static images rather than a moving
23 animation."

24 A common example of the problem is characterized by a status ribbon
25 indicator that is displayed while a system is processing information. Usually, the

1 indicator progresses, for instance, in a left-to-right direction that indicates that
2 some processing is going on. However, there are times when the status indicator
3 stops and the user cannot tell if the system has locked up or is still processing
4 something. This is a result of the interactive portion of the animator not providing
5 updated data to the display portion in time to meet the frame rate refresh.

6 Such a system is particularly poor for videos. With a video, a frame can be
7 dropped - resulting in a glitch, or artifact - simply because the system was busy
8 taking care of another task (that could be as minor as moving a mouse cursor).

9 In a factorized animation system, the low-level animation subsystem 148
10 runs asynchronously from the high-level animation subsystem 122, so it continues
11 to run (i.e., draw) even when the high-level animation subsystem 122 gets tied up
12 with a task that prevents it from re-drawing a frame before the frame is refreshed.
13 The factorized animation system, therefore, seamlessly displays an animation
14 operation and, therefore, provides a superior user experience.

15 **Exemplary Methodological Implementation: Basic Operation**

16 **Fig. 3** is a flow diagram that depicts basic operation of a factorized
17 animation/rendering system, such as shown in Fig. 1 and Fig. 2. In the following
18 discussion, continuing reference will be made to elements and reference numerals
19 shown in Figures 1 and 2.

20 In a factorized animation system (i.e., an animation system that is factored
21 into at least two distinct animation subsystems that function asynchronously with
22 each other), a user (application) tells a high-level system, through system APIs,
23 how graphics are supposed to appear on a display. The high-level system uses a
24 specialized communication protocol to direct a low-level system to create a
25 structure similar to one set up in the high-level system.

1 Since the high-level system is configured to create several elements in the
2 low-level system, it necessarily follows that the communication protocol -
3 discussed in greater detail, below, with respect to Fig. 4 - includes several “create”
4 messages or commands. The flow diagram shown in Fig. 3 covers some basic
5 steps that occur when the high-level animation subsystem 122 communicates with
6 the low-level animation subsystem 148 to set up structure therein.

7 At block 300, the high-level animation subsystem 122 sends a message to
8 create an animation object 220 in the low-level animation subsystem 148. The
9 high-level animation subsystem 122 then communicates with the low-level
10 animation subsystem 148 to create a clock 224 (i.e., a timeline) to include with the
11 animation object 220 that was just created (block 302).

12 It is noted that there is not necessarily a 1:1 ratio between animation objects
13 and timelines. In other words, one timeline may control more than one animation
14 object. Therefore, in the step described above for block 302, the timeline may not
15 need to be created if one already exists and if the newly-created animation object
16 can be associated with the existing timeline.

17 At this point, the application 120 may create another animation object or
18 modify an existing animation. If the application 120 is configured to create
19 another animation object (“animation object” branch, block 304), then the high-
20 level animation subsystem 122 sends a communication protocol message to the
21 low-level animation subsystem 138 to create another animation object 220 at
22 block 300. The high-level animation subsystem 122 then sends a message at block
23 302 to associate a timeline with the newly-created animation object 220. The
24 timeline may be created or it may be an existing timeline.

1 The application 120 is also configured to support modifications. The high-
2 level animation subsystem 122 can send a message to modify an animation object
3 220 and/or a timeline associated therewith (“modification” branch, block 304). In
4 addition to a straightforward modification that modifies an aspect of an animation,
5 a “pause” operation and a “stop” operation are also modifications. A “pause”
6 operation can be coded in terms of a typical modification, e.g., “at time 0, clock is
7 10; and at time 10, clock is 10.” A “stop” operation is initiated with a “remove
8 animation” message.

9 If the modification is a “remove animation” message (“Yes” branch, block
10 306, then the animation is halted at block 308. Otherwise (“No” branch, block
11 306), the high-level animation system 122 sends one or more messages configured
12 to modify an existing animation (block 310). Modifying an existing animation
13 may also mean modifying the timeline associated with an existing animation.

14 As previously mentioned, the above flow diagram depicts a general
15 methodology for configuring a low-level animation subsystem 148 by sending
16 communication protocol messages from a high-level animation subsystem 122. At
17 least one more specific methodological implementation utilizing the
18 communication protocol will be discussed in greater detail below, after the
19 specifics of the communication protocol are discussed.

20 **Communication Protocol**

21 Figures 4a, 4b and 4c are tables outlining messages/commands of a
22 communication protocol for use in governing transmissions between a high-level
23 animation system and a low-level animation system. In the following discussion,
24 continuing reference will be made to the elements and reference numerals shown
25 and described in Fig. 1 and Fig. 2.

1 **Fig. 4a** depicts a table 400 that describes messages sent from the high-level
2 timing engine 212 to the low-level timing engine 222. A “Create Clock” message
3 402 causes the low-level timing engine 222 to create a clock for an animation.
4 The “Create Clock” message 402 includes “Create Clock Parameters” 404, which
5 are the initial clock properties. Clock properties may include, but are not limited
6 to, duration (from a begin time), parent clock (to which all times in this clock are
7 relative), speed (relative to its parent), acceleration and deceleration. The
8 acceleration and deceleration parameters specify the percentage of the “duration”
9 time that is spent “speeding up” or “slowing down” the animation.

10 An “Update Properties” message 406 provides an update to an existing
11 clock’s properties. “Update Properties Parameters” 408 include the target clock
12 (i.e. the clock having the properties to be updated), and the updated properties and
13 values. An “Add Interval” message 410 instructs the low-level timing engine to
14 add an interval to an existing animation (object) and includes “Add Interval
15 Parameters” 412 that identify a target clock and properties of the interval that is
16 being added to the target clock.

17 The protocol includes a “Reset Synchronization Slip” message 414 and an
18 associated “Reset Synchronization Slip Parameter” 416 that are used to maintain
19 synchronization between the high-level animation subsystem 122 and the low-
20 level animation subsystem 148. How this is implemented is a matter of a
21 particular system design.

22 In at least one implementation, the synchronization process is described by
23 the following example. If the low-level animation subsystem 148 is controlling a
24 video file and it detects that the video is falling behind (due to, for example,
25 network traffic), then the low-level animation subsystem 148 sends a

1 “Synchronize With Media Slip” message and parameter(s) (see Fig. 4b, below) to
2 the high-level animation subsystem 122 and stores the fact that a slip has occurred
3 and the magnitude of the slip.

4 When the high-level animation subsystem 122 (specifically, the high-level
5 timing engine 212) finishes updating the intervals to take the slip into account, it
6 sends the “Reset Synchronization Slip” message 414 and an associated “Reset
7 Synchronization Slip Parameter” 416 to tell the low-level animation subsystem
8 148 to reset the slip offset to zero.

9 For further information on this particular protocol message, please refer to
10 Fig. 4b, below.

11 A “Remove All Intervals” message 418 may be sent to the low-level timing
12 engine 222 to remove any existing intervals associated with a particular clock.
13 The target clock is identified in a “Remove All Intervals Parameter” 420. A
14 “Delete Clock” message 422 is also included in the communication protocol and is
15 sent to the low-level timing engine 222 to remove a target clock - identified in a
16 “Delete Clock Parameter” 424.

17 **Fig. 4b** depicts a table 430 that describes a message in the communication
18 protocol that is sent from the low-level timing engine 222 to the high-level timing
19 engine 212. A “Synchronize With Media Slip” message 432 provides the high-
20 level timing engine 212 with an amount that the high-level timing engine 212
21 should slip a particular clock to sync-up with a media component. *** NOTE: Is
22 this even close to correct? *** “Synchronize With Media Slip Parameters” 434
23 include a target clock that is the clock to slip, and a slip amount that identifies the
24 magnitude of the slip.

1 The message sends data back to the high-level animation subsystem 122
2 from the low-level animation subsystem. This is necessary due to the existence of
3 a latency in the communication channel 204. One instance in which this may
4 happen is when a user activates a “pause” command. Since the systems
5 communicate with each other asynchronously, the low-level animation subsystem
6 may process a number of frames before the high-level animation subsystem can
7 catch up to the low-level animation subsystem. As a result, the low-level system
8 runs long.

9 But this is taken care of with the “Synchronize With Media Slip” message
10 432 that gets information to the high-level animation subsystem 122 that the clock
11 associated with the animation needs to “slip” or “synch up” with the low-level
12 clock 224 to get the subsystems back in synchronization.

13 **Fig. 4c** depicts a table 440 that describes messages sent from the high-level
14 animation objects 210 to the low-level animation objects 220. The table 440
15 includes a “Create Animation” message 442 that commands the low-level
16 animation subsystem 148 to create a new animation object. “Create Animation
17 Parameters” 444 include an output value type that identifies the type of the
18 property being animated, an animation function, an animation function and a
19 controlling clock.

20 For example, if the angle of a rotation transform is animated, then the
21 output type is “double-precision floating point number.” If one of the end points
22 of a line is animated, then the output type is “point.” If the color of a rectangle is
23 animated, then the output type is “color.”

24 Animation functions are known in the art and any such animation function
25 may be designated here. An example of a simple animation function is “from 5 to

1 10.” This means that at the beginning of the interval, the animation function
2 outputs a value of 5, at the end it outputs a 10, and at halfway through the
3 animation function outputs a value of 7.5, etc. A more complex animation
4 function is “from 5 to 6 to 10.” Here, the output value halfway through the
5 animation is 6. The function can be made more complex by specifying not only a
6 list of values, but a list of times at which the function should output those values
7 (with everything in between being interpolated).

8 It is noted that the animation functions described above are exemplary only
9 and are not intended to limit the scope of the appended claims to these animation
10 functions. Any computer operation that can be considered to be an “animation
11 function” may be used in this context.

12 Finally, as previously discussed, the controlling clock for a new animation
13 object may be a new clock or it may be an existing clock that may or may not be
14 controlling one or more other animation objects.

15 The table 440 also includes an “Update Animation” message 446 that
16 provides information to update an existing animation in the low-level animation
17 subsystem 148. “Update Animation Parameters” 448 include a target animation
18 (i.e. the animation to be updated), and updated properties (properties to be updated
19 and updated values for the properties).

20 A “Create Animation Collection” message 450 identifies multiple
21 animations in the low-level animation objects [database] 220 that are to be
22 grouped as a collection. A “Create Animation Collection Parameter” 452
23 identifies an initial list of animations that are to be grouped as a collection.

24 Animation collections are also well known in the art. Animation
25 collections exist to allow a developer to create complex animation behaviors that

1 cannot be expressed with a single linear interpolation. For example, one
2 animation can move a point up and down repeatedly, while a second animation can
3 move a point to the right. When both animations are run together, the point
4 appears to follow a wave pattern to the right. Animation collections can be edited,
5 hence the existence of the messages in the protocol described below.

6 An “Add Animation To Collection” message 454 tells the low-level system
7 to add an animation identified in an “Add Animation To Collection Parameter”
8 456 to an animation collection that is also identified in the “Add Animation To
9 Collection Parameter” 456. Similarly, a “Remove Animation From Collection”
10 message 458 is included with a “Remove Animation From Collection Parameter”
11 460 that identifies a target animation collection and an animation in the collection
12 that is to be removed from the collection.

13 The table 440 also includes a “Create Static Value” message 462. A static
14 value is used in the case where the low-level animation subsystem 148 cannot
15 support an animation provided by the application 120. This situation is described
16 in greater detail, below, with respect to Fig. 6. Associated with the “Create Static
17 Value” message 462 is a “Create Static Value Parameter” 464 that identifies a
18 value type and an initial value for the static value that is created.

19 An “Update Static Value” message 466 provides an update to a static value.
20 An “Update Static Value Parameter” 468 identifies a target static value object and
21 a new value for the static value object so identified.

22 The messages and parameters shown in Fig. 4(a - c) provide a streamlined,
23 efficient way in which animations rendered by a low-level animation subsystem
24 148 can be asynchronously controlled by a high-level animation subsystem. Since
25 one or more messages do not have to be sent in every rendering refresh frame, the

1 display can render frames at a high, constant rate. At the same time, the high-level
2 system can use processing time as it becomes available to provide data to the low-
3 level system, i.e. at a slower and varied refresh rate.

4 **Exemplary Methodological Implementation: Fully Connected System**

5 **Fig. 5** is a flow diagram depicting an exemplary methodological
6 implementation of an animation process in a rendering system having a fully
7 connected high level animation system and low level animation system.

8 The scenario depicted in Fig. 5 is: An application developer wants to draw
9 an animated straight line going from a fixed anchor point to a second point that
10 moves over time. The developer needs to provide three things: (1) the rendering
11 operation (e.g. draw a line between static point A and animate point B); (2) the
12 animation function (e.g. point B goes from (0,0) to (100, 100); and (3) the timing
13 of the animation (e.g. start in five seconds and run for ten seconds, then repeat
14 once going backwards).

15 Fig. 5 shows the response of the high-level animation subsystem 122 to the
16 application 120 containing the information discussed above. At block 500, the
17 high-level animation subsystem 122 receives the information from the application
18 120. The high-level animation subsystem 122 creates a rendering object at block
19 502, the rendering object representing a line. At block 504, the high-level
20 animation subsystem 122 creates an animation function that outputs a value from
21 (0,0) to (100,100). At block 506, the high-level animation subsystem 122 creates a
22 clock that starts in five seconds, runs for ten second and then runs backwards.

23 The low-level animation subsystem 148 can support these objects, so the
24 high-level animation subsystem 120 sends messages to the low-level animation
25

1 subsystem 148 (block 508) that generates traffic on the communication channel
2 204 that is depicted in blocks 510 - 518.

3 Block 510: **Create Clock 1**; Set parameters for clock 1, duration = 10.

4 Block 512: **Add Interval** for Clock 1; From “now+5” to “now+15”, go
5 from t=0 to t=10.

6 Block 514: **Add Interval** for Clock 1; From “now+15” to “now+25”, go
7 from t=10 down to t=1.

8 Block 516: **Create Animation 1**; Set parameters for animation 1, from
9 (0,0) to (100,100), clock=clock 1.

10 Block 518: Create [DrawLine] instruction referencing animation 1.

11 In this example, the DrawLine function is representative of any drawing
12 instruction supported by a protocol for rendering animations. The DrawLine
13 instruction is exemplary only.

14 After the messages have been transmitted from the high-level animation
15 subsystem 122 to the low-level animation subsystem 148, the low-level animation
16 subsystem 148 runs independently of the high-level subsystem 148, and updates
17 the position of the line at the highest possible frame rate. No further messages
18 need to be sent between the two subsystems 122, 148.

19 As has been shown, the communication protocol introduced herein
20 optimizes the communication channel 204 by simplifying the amount of data that
21 is put into the communication channel 204. If the short messages and parameters
22 discussed above were not used, a large number of operations would have to be sent
23 through the pipeline to synchronize the systems.

24 But, in essence, what is put into the communication channel 204 is a list of
25 intervals for a number of clocks. For a clock an interval is defined, say, as

1 'between time 0 and time 10, the clock goes from 0 to 10' (this defines a real-
2 world clock). Or, maybe the clock is desired to run twice as slow. In this case,
3 then 'between time 0 and time 10, the clock goes from 0 to 5'.

4 By defining a list of timing intervals, every operation on a clock can be
5 defined. For example, a pause operation can be defined with a clock interval: 'at
6 time 0, the clock is 10 and at time 10, the clock is 10'. Therefore, every operation
7 can be expressed as a linear interval.

8 9 **Exemplary Methodological Implementation: Partly Connected System**

10 **Fig. 6** is a flow diagram depicting an exemplary methodological
11 implementation of an animation process in a rendering system having a partially
12 connected high level animation system and low level animation system.

13 The scenario for this example is: An application developer wants to draw an
14 animated straight line going from a fixed anchor point to a second point that
15 moves over time on a custom path. The developer needs to provide the following:
16 (1) the rendering operation (e.g. draw a line between static point A and animate
17 point B); (2) the custom animation function (e.g. call custom animation 1 to
18 compute the position of point B); and (3) the timing of the animation (e.g. start in
19 five seconds and run for ten seconds, then repeat once going backwards).

20 The high-level animation subsystem 122 receives the application
21 information at block 600. In response to the application's request, the high-level
22 animation subsystem 122 creates a rendering object (block 602) representing the
23 line, an animation function (block 604) that is configured to call back user code
24 202, and a clock (block 606) that starts in five seconds, runs for ten seconds and
25

1 then runs backwards. The appropriate messages are sent to the low-level
2 animation subsystem 148 at block 608.

3 Since the custom animation requires user code, it can't be processed by the
4 low-level subsystem, so this generates the following initial traffic on the
5 communication channel:

6 **Block 610: Create Static Point Value 1.**

7 **Block 612: Create [DrawLine] Instructions Referencing Point Value 1.**
8 (DrawLine instruction is exemplary only; see above).

9 At block 614, the low-level animation subsystem 148 receives the
10 transmitted data from the high-level animation subsystem 122. The low-level
11 animation subsystem 148 treats the animation as a static line.

12 On every frame after the initial transmission of data, the high-level
13 animation subsystem 122 calls the user code 202 to compute the position of point
14 B (block 620) and receives the new position value at block 622. The high-level
15 animation subsystem 122 then sends the following message over the
16 communication channel 204 to the low-level animation subsystem to update the
17 value of the static point value 1:

18 **Block 626: Update Static Point Value with New Value for Point B.**

19 **EXEMPLARY COMPUTER ENVIRONMENT**

20 The various components and functionality described herein are
21 implemented with a computing system. Fig. 7 shows components of typical
22 example of such a computing system, i.e. a computer, referred by to reference
23 numeral 700. The components shown in Fig. 7 are only examples, and are not
24 intended to suggest any limitation as to the scope of the functionality of the
25

1 invention; the invention is not necessarily dependent on the features shown in Fig.
2 7.

3 Generally, various different general purpose or special purpose computing
4 system configurations can be used. Examples of well known computing systems,
5 environments, and/or configurations that may be suitable for use with the
6 invention include, but are not limited to, personal computers, server computers,
7 hand-held or laptop devices, multiprocessor systems, microprocessor-based
8 systems, set top boxes, programmable consumer electronics, network PCs,
9 minicomputers, mainframe computers, distributed computing environments that
10 include any of the above systems or devices, and the like.

11 The functionality of the computers is embodied in many cases by computer-
12 executable instructions, such as program modules, that are executed by the
13 computers. Generally, program modules include routines, programs, objects,
14 components, data structures, etc. that perform particular tasks or implement
15 particular abstract data types. Tasks might also be performed by remote
16 processing devices that are linked through a communications network. In a
17 distributed computing environment, program modules may be located in both local
18 and remote computer storage media.

19 The instructions and/or program modules are stored at different times in the
20 various computer-readable media that are either part of the computer or that can be
21 read by the computer. Programs are typically distributed, for example, on floppy
22 disks, CD-ROMs, DVD, or some form of communication media such as a
23 modulated signal. From there, they are installed or loaded into the secondary
24 memory of a computer. At execution, they are loaded at least partially into the
25 computer's primary electronic memory. The invention described herein includes

1 these and other various types of computer-readable media when such media
2 contain instructions programs, and/or modules for implementing the steps
3 described below in conjunction with a microprocessor or other data processors.
4 The invention also includes the computer itself when programmed according to
5 the methods and techniques described below.

6 For purposes of illustration, programs and other executable program
7 components such as the operating system are illustrated herein as discrete blocks,
8 although it is recognized that such programs and components reside at various
9 times in different storage components of the computer, and are executed by the
10 data processor(s) of the computer.

11 With reference to Fig. 7, the components of computer 700 may include, but
12 are not limited to, a processing unit 702, a system memory 704, and a system bus
13 706 that couples various system components including the system memory to the
14 processing unit 702. The system bus 706 may be any of several types of bus
15 structures including a memory bus or memory controller, a peripheral bus, and a
16 local bus using any of a variety of bus architectures. By way of example, and not
17 limitation, such architectures include Industry Standard Architecture (ISA) bus,
18 Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video
19 Electronics Standards Association (VESA) local bus, and Peripheral Component
20 Interconnect (PCI) bus also known as the Mezzanine bus.

21 Computer 700 typically includes a variety of computer-readable media.
22 Computer-readable media can be any available media that can be accessed by
23 computer 700 and includes both volatile and nonvolatile media, removable and
24 non-removable media. By way of example, and not limitation, computer-readable
25 media may comprise computer storage media and communication media.

1 “Computer storage media” includes volatile and nonvolatile, removable and non-
2 removable media implemented in any method or technology for storage of
3 information such as computer-readable instructions, data structures, program
4 modules, or other data. Computer storage media includes, but is not limited to,
5 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
6 digital versatile disks (DVD) or other optical disk storage, magnetic cassettes,
7 magnetic tape, magnetic disk storage or other magnetic storage devices, or any
8 other medium which can be used to store the desired information and which can be
9 accessed by computer 700. Communication media typically embodies computer-
10 readable instructions, data structures, program modules or other data in a
11 modulated data signal such as a carrier wave or other transport mechanism and
12 includes any information delivery media. The term “modulated data signal”
13 means a signal that has one or more of its characteristics set or changed in such a
14 manner as to encode information in the signal. By way of example, and not
15 limitation, communication media includes wired media such as a wired network or
16 direct-wired connection and wireless media such as acoustic, RF, infrared and
17 other wireless media. Combinations of any of the above should also be included
18 within the scope of computer readable media.

19 The system memory 704 includes computer storage media in the form of
20 volatile and/or nonvolatile memory such as read only memory (ROM) 708 and
21 random access memory (RAM) 710. A basic input/output system 712 (BIOS),
22 containing the basic routines that help to transfer information between elements
23 within computer 700, such as during start-up, is typically stored in ROM 708.
24 RAM 710 typically contains data and/or program modules that are immediately
25 accessible to and/or presently being operated on by processing unit 702. By way

1 of example, and not limitation, Fig. 7 illustrates operating system 714, application
2 programs 716, other program modules 718, and program data 720.

3 The computer 700 may also include other removable/non-removable,
4 volatile/nonvolatile computer storage media. By way of example only, Fig. 7
5 illustrates a hard disk drive 722 that reads from or writes to non-removable,
6 nonvolatile magnetic media, a magnetic disk drive 724 that reads from or writes to
7 a removable, nonvolatile magnetic disk 726, and an optical disk drive 728 that
8 reads from or writes to a removable, nonvolatile optical disk 730 such as a CD
9 ROM or other optical media. Other removable/non-removable,
10 volatile/nonvolatile computer storage media that can be used in the exemplary
11 operating environment include, but are not limited to, magnetic tape cassettes,
12 flash memory cards, digital versatile disks, digital video tape, solid state RAM,
13 solid state ROM, and the like. The hard disk drive 722 is typically connected to
14 the system bus 706 through a non-removable memory interface such as data media
15 interface 732, and magnetic disk drive 724 and optical disk drive 728 are typically
16 connected to the system bus 706 by a removable memory interface such as
17 interface 734.

18 The drives and their associated computer storage media discussed above
19 and illustrated in Fig. 7 provide storage of computer-readable instructions, data
20 structures, program modules, and other data for computer 700. In Fig. 7, for
21 example, hard disk drive 722 is illustrated as storing operating system 715,
22 application programs 717, other program modules 719, and program data 721.
23 Note that these components can either be the same as or different from operating
24 system 714, application programs 716, other program modules 718, and program
25 data 720. Operating system 715, application programs 717, other program

1 modules 719, and program data 721 are given different numbers here to illustrate
2 that, at a minimum, they are different copies. A user may enter commands and
3 information into the computer 700 through input devices such as a keyboard 736
4 and pointing device 738, commonly referred to as a mouse, trackball, or touch
5 pad. Other input devices (not shown) may include a microphone, joystick, game
6 pad, satellite dish, scanner, or the like. These and other input devices are often
7 connected to the processing unit 702 through an input/output (I/O) interface 740
8 that is coupled to the system bus, but may be connected by other interface and bus
9 structures, such as a parallel port, game port, or a universal serial bus (USB). A
10 monitor 742 or other type of display device is also connected to the system bus
11 706 via an interface, such as a video adapter 744. In addition to the monitor 742,
12 computers may also include other peripheral output devices 746 (e.g., speakers)
13 and one or more printers 748, which may be connected through the I/O interface
14 740.

15 The computer may operate in a networked environment using logical
16 connections to one or more remote computers, such as a remote computing device
17 750. The remote computing device 750 may be a personal computer, a server, a
18 router, a network PC, a peer device or other common network node, and typically
19 includes many or all of the elements described above relative to computer 700.
20 The logical connections depicted in Fig. 7 include a local area network (LAN) 752
21 and a wide area network (WAN) 754. Although the WAN 754 shown in Fig. 7 is
22 the Internet, the WAN 754 may also include other networks. Such networking
23 environments are commonplace in offices, enterprise-wide computer networks,
24 intranets, and the like.

1 When used in a LAN networking environment, the computer 700 is
2 connected to the LAN 752 through a network interface or adapter 756. When used
3 in a WAN networking environment, the computer 700 typically includes a modem
4 758 or other means for establishing communications over the Internet 754. The
5 modem 758, which may be internal or external, may be connected to the system
6 bus 706 via the I/O interface 740, or other appropriate mechanism. In a networked
7 environment, program modules depicted relative to the computer 700, or portions
8 thereof, may be stored in the remote computing device 750. By way of example,
9 and not limitation, Fig. 7 illustrates remote application programs 760 as residing
10 on remote computing device 750. It will be appreciated that the network
11 connections shown are exemplary and other means of establishing a
12 communications link between the computers may be used.

13 14 **Conclusion**

15 The communication protocol for the systems and methods as described thus
16 provide a way to synchronize a high-level animation system with a low-level
17 animation system so that the low-level system can run at a fast, constant frame
18 refresh rate, while the high-level system can run at a variable frame rate that is
19 optimized for interactivity. Superior graphics can be achieved without sacrificing
20 any interaction capability. The protocol also minimizes the amount of data
21 transmitted between the two animation systems, thus saving overhead and
22 optimizing performance.

23 Although details of specific implementations and embodiments are
24 described above, such details are intended to satisfy statutory disclosure
25 obligations rather than to limit the scope of the following claims. Thus, the

1 invention as defined by the claims is not limited to the specific features described
2 above. Rather, the invention is claimed in any of its forms or modifications that
3 fall within the proper scope of the appended claims, appropriately interpreted in
4 accordance with the doctrine of equivalents.

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25